

CLAIMS

What is claimed is:

- 1 1. A method of bandwidth management for use in a packet forwarding device coupled
2 to a network, the method comprising the steps of:
3 receiving at a packet forwarding device information indicative of one or more traffic
4 groups;
5 receiving at the packet forwarding device one or more bandwidth parameters for at
6 least one of the one or more traffic groups;
7 receiving at a first port of a plurality of ports a packet associated with the at least
8 one traffic group; and
9 scheduling the packet for transmission from a second port of the plurality of ports
10 based upon the one or more bandwidth parameters for the at least one traffic
11 group with which the packet is associated.
- 1 2. The method of claim 1, wherein the network employs a non-deterministic access
2 protocol.
- 1 3. The method of claim 2, wherein the non-deterministic access protocol is Carrier
2 Sense Multiple Access with Collision Detection (CSMA/CD).
- 1 4. The method of claim 1, wherein the one or more bandwidth parameters include a
2 minimum bandwidth.
- 1 5. The method of claim 1, wherein the one or more bandwidth parameters include a
2 maximum bandwidth.

- 1 6. The method of claim 1, wherein the one or more bandwidth parameters include a
2 peak bandwidth.
- 1 7. The method of claim 1, wherein the one or more bandwidth parameters include a
2 maximum delay.
- 1 8. The method of claim 1, wherein the one or more bandwidth parameters include a
2 relative priority.
- 1 9. The method of claim 1, further comprising the steps of:
2 classifying the packet as being associated with the at least one traffic group; and
3 determining a quality of service queue with which the at least one traffic group is
4 associated.
- 1 10. The method of claim 1, further comprising the steps of:
2 enqueueing the packet onto a queue associated with the traffic group;
3 determining a current bandwidth metric for the queue; and
4 dequeuing the packet from the queue if the current bandwidth metric meets a
5 predetermined relationship with the one or more bandwidth parameters.
- 1 11. The method of claim 10, wherein the current bandwidth metric is evaluated
2 periodically at the expiration of a predetermined time period, and wherein the step
3 of determining a current bandwidth metric for the queue further comprises the steps
4 of:
5 determining an actual bandwidth for a prior time period;
6 determining a bandwidth metric for the prior time period; and

combining a portion of the actual bandwidth for the prior time period with a portion of the bandwidth metric for the prior time period to arrive at the current bandwidth metric.

12. A method of bandwidth management and traffic prioritization for use in a network of devices, the method comprising the steps of:
defining at a packet forwarding device information indicative of one or more traffic groups;
defining at the packet forwarding device information indicative of a quality of service (QoS) policy for at least one of the one or more traffic groups;
receiving a packet at a first port of a plurality of ports;
identifying a first traffic group of the one or more traffic groups with which the packet is associated; and
scheduling the packet for transmission from a second port of the plurality of ports based upon the QoS policy for the first traffic group, and wherein the scheduling is independent of end-to-end signaling.

13. The method of claim 12, wherein the network of devices employs a non-deterministic access protocol.

14. The method of claim 13, wherein the non-deterministic access protocol is Carrier Sense Multiple Access with Collision Detection (CSMA/CD).

15. The method of claim 12, further comprising the steps of:
providing a plurality of QoS queues; and
mapping the first traffic group to a first QoS queue of the plurality of QoS queues.

- 1 16. The method of claim 15, further comprising the steps of:
2 determining a current bandwidth metric for each of the plurality of QoS queues;
3 dividing the plurality of QoS queues into at least a first group and a second group
4 based upon the current bandwidth metrics and a minimum bandwidth
5 requirement associated with each of the plurality of QoS queues;
6 if the first group includes at least one QoS queue, then transmitting a packet from
7 the at least one QoS queue; otherwise transmitting a packet from a QoS
8 queue associated with the second group.
- 1 17. A method of bandwidth management and traffic prioritization for use in a network
2 of devices, the method comprising the steps of:
3 receiving at a packet forwarding device information indicative of one or more traffic
4 groups;
5 receiving at the packet forwarding device information defining a quality of service
6 (QoS) policy for at least one of the one or more traffic groups, the QoS
7 policy including at least a minimum bandwidth;
8 providing a plurality of queues at each of a plurality of output ports;
9 associating the one or more traffic groups with the plurality of queues based upon
10 the minimum bandwidth; and
11 scheduling a packet for transmission from one of the plurality of queues onto the
12 network.
- 1 18. The method of claim 17, wherein the information indicative of the one or more
2 traffic groups includes Internet Protocol (IP) subnet membership.

- 1 19. The method of claim 18, wherein the information indicative of the one or more
2 traffic groups includes a media access control (MAC) address.
- 1 20. The method of claim 17, wherein the information indicative of the one or more
2 traffic groups includes a virtual local area network (VLAN) identifier.
- 1 21. A method of bandwidth management and traffic prioritization for use in a network
2 of devices, the method comprising the steps of:
3 providing a plurality of quality of service (QoS) queues at each of a plurality of
4 output ports, each of the plurality of QoS queues associated with a
5 minimum queue bandwidth requirement;
6 adding a packet to one of the plurality of QoS queues based upon a traffic group
7 with which the packet is associated; and
8 scheduling a next packet for transmission onto the network from one of the plurality
9 of QoS queues at a particular output port of the plurality of output ports by:
10 determining a current bandwidth metric for each of the plurality of QoS
11 queues,
12 dividing the plurality of QoS queues into at least a first group and a second
13 group based upon the current bandwidth metrics and the minimum
14 queue bandwidth requirements, and
15 if at least one QoS queue of the plurality of QoS queues, so divided, is
16 associated with the first group, then transmitting a packet from the at
17 least one QoS queue; otherwise transmitting a packet from a QoS
18 queue of the plurality of QoS queues associated with the second
19 group.

1 22. The method of claim 21, wherein the current bandwidth for a particular QoS queue
2 is calculated as follows:

3
$$\text{CURR_BW}_i = W1 \times \text{CURR_BW}_i + W2 \times \text{ACT_BW}_i;$$

4 where:

5 CURR_BW_i represents the current bandwidth for a particular QoS queue,

6 $W1$ represents a first weighting factor,

7 $W2$ represents a second weighting factor, and

8 ACT_BW_i represents the actual bandwidth received by the particular QoS queue in
9 a previous time interval.

1 23. The method of claim 22, wherein $W1 = (W-1)/W$,

2 $W2 = 1/W$, and the previous time interval is the most recent time interval.

1 24. The method of claim 21, further comprising the step of selecting among QoS
2 queues in the same group based upon relative queue priorities associated with the
3 QoS queues.

1 25. The method of claim 21, further comprising the step of selecting among QoS
2 queues in the same group based upon a round robin selection scheme.

1 26. The method of claim 21, further comprising the step of selecting among QoS
2 queues in the same group based upon a least recently used (LRU) selection scheme.

1 27. The method of claim 21, wherein the first group comprises QoS queues associated
2 with a minimum queue bandwidth requirement that is less than the corresponding
3 QoS queue's current bandwidth metric, and wherein the second group comprises

QoS queues associated with a minimum queue bandwidth requirement that is greater than or equal to the corresponding QoS queue's current bandwidth metric.

28. A method of bandwidth management for use in a packet forwarding device participating in a connectionless network, the method comprising the steps of: receiving at a packet forwarding device information indicative of one or more traffic groups; receiving at the packet forwarding device one or more bandwidth parameters for at least one of the one or more traffic groups; receiving at a first port of a plurality of ports a packet associated with the at least one traffic group; and scheduling the packet for transmission from a second port of the plurality of ports based upon the one or more bandwidth parameters for the traffic group with which the packet is associated.

29. A packet forwarding device for use in a network employing a non-deterministic assess protocol, the packet forwarding device comprising: an input unit configured to receive information defining one or more traffic groups and one or more associated bandwidth parameters; a plurality of ports configured to transmit packets onto an attached network segment, each port having a plurality of queues and configured to select a queue of the plurality of queue from which to transmit a next packet based upon the one or more bandwidth parameters; and a filtering and forwarding engine coupled to the plurality of ports and configured to process received packets, the filtering and forwarding engine identifying a traffic group of the one or more traffic groups with which a received packet

12 is associated and queuing the received packet for transmission from one of
13 the plurality of ports based upon the identified traffic group.

1 30. A packet forwarding device for use in a network employing a non-deterministic
2 assess protocol, the packet forwarding device comprising:
3 a filtering and forwarding engine configured to forward received packets based
4 upon a traffic group with which the packet is associated; and
5 a plurality of ports coupled to the filtering and forwarding engine, each port of the
6 plurality of ports configured to receive packets from the filtering and
7 forwarding engine, each port of the plurality of ports having a plurality of
8 Quality of Service (QoS) queues associated with a minimum queue
9 bandwidth requirement, each port of the plurality of ports further configured
10 to schedule a packet for transmission onto the network by
11 determining a current bandwidth metric for each of the plurality of QoS
12 queues,
13 dividing the plurality of QoS queues into at least a first group and a second
14 group based upon the current bandwidth metrics and the minimum
15 queue bandwidth requirements, and
16 if at least one QoS queue of the plurality of QoS queues, so divided, is
17 associated with the first group, then transmitting a packet from the at
18 least one QoS queue; otherwise transmitting a packet from a QoS
19 queue of the plurality of QoS queues associated with the second
20 group.

- 1 31. The packet forwarding device of claim 30, wherein the plurality of ports are further
2 configured to select among QoS queues in the same group based upon relative
3 queue priorities associated with the QoS queues.
- 1 32. The packet forwarding device of claim 30, wherein the plurality of ports are further
2 configured to select among QoS queues in the same group based upon a round
3 robin selection scheme.
- 1 33. The packet forwarding device of claim 30, wherein the plurality of ports are further
2 configured to select among QoS queues in the same group based upon a least
3 recently used (LRU) selection scheme.
- 1 34. The packet forwarding device of claim 30, wherein the first group comprises QoS
2 queues associated with a minimum queue bandwidth requirement that is less than
3 the corresponding QoS queue's current bandwidth metric, and wherein the second
4 group comprises QoS queues associated with a minimum queue bandwidth
5 requirement that is greater than or equal to the corresponding QoS queue's current
6 bandwidth metric.
- 1 35. A machine-readable medium having stored thereon data representing sequences of
2 instructions, said sequences of instructions which, when executed by a processor,
3 cause said processor to perform the steps of:
4 receiving at a packet forwarding device information indicative of one or more traffic
5 groups;
6 receiving at the packet forwarding device information defining a quality of service
7 (QoS) policy for at least one of the one or more traffic groups;

8 receiving a packet at a first port of a plurality of ports;
9 identifying a first traffic group of the one or more traffic groups with which the
10 packet is associated; and
11 scheduling the packet for transmission from a second port of the plurality of ports
12 based upon the QoS policy for the first traffic group, and wherein the
13 scheduling is independent of end-to-end signaling.